

REMARKS

The Examiner is thanked for the performance of a thorough search. By this amendment, Claims 5, 10-15, 24 and 29 have been canceled, Claims 1, 16 and 20 amended and new Claims 30-32 added. Hence, Claims 1-4, 6-9, 16-23, 25-28 and 30-32 are pending in this application. The amendments to the claims and new claims do not add any new matter to this application. Claims 1 and 16 have been amended to include the limitations of dependent Claim 5. Claim 20 has been amended to include the limitations of dependent Claim 24. All issues raised in the Office Action mailed September 26, 2002 are addressed hereinafter.

REJECTION OF CLAIMS 1-29 UNDER 35 U.S.C. §103(a)

Claims 1-29 were rejected under 35 U.S.C. §103(a) as being unpatentable over *Abdelnur et al.*, U.S. Patent No. 6,212,640 (hereinafter "*Abdelnur*") in view of *Smith et al.*, U.S. Patent No. 5,835,724 (hereinafter "*Smith*"). This rejection is now moot with respect to canceled Claims 5, 10-15, 24 and 29. It is respectfully submitted that Claims 1-4, 6-9, 16-23 and 25-28, as amended, are patentable over *Abdelnur* and *Smith*, alone or in combination, for at least the reasons provided hereinafter.

CLAIM 1

Claim 1 recites a method for processing data on a distributed computing system that includes a plurality of nodes that requires the steps of:

"maintaining mapping data that specifies work that can be performed by each of the plurality of nodes;
in response to receiving a first work request to perform first work from a first process on a first node from the plurality of nodes, determining based upon the first work and the mapping data, that the first work is to be performed on a second node from the plurality of nodes;

providing the first work request to a second process on the second node, wherein the first work request specifies that the first process is to directly receive results of the first work;
determining based upon the first work and the mapping data, that the first work is also to be performed on a third node from the plurality of nodes, and
providing a second work request to a third process on the third node, wherein the second work request specifies that results of the first work performed on the third node are to be provided directly to the first process.

It is respectfully submitted that Claim 1 requires at least several limitations that are not in any way taught or suggested by *Abdelnur* and *Smith*, alone or in combination.

Abdelnur discloses an approach for sharing resources in a network environment. In a first scenario, a client application 410 is a trusted application to server 460. In this situation, client application 410 establishes a communication link 440 with server 460 and submits a request directly to server 460. Server 460 processes the request and provides a response directly back to client application 410.

In a second scenario, client application 410 is not a trusted application to server 460 and therefore cannot directly submit requests to server 460, or directly receive a response from server 460. Client application 410 is, however, a trusted application to Web server 480. In this situation, client application 410 submits its request to Web server 480. The request is processed by a servlet 490 executing on Web server 480. Servlet 490 submits the request to server 460 on behalf of client application 410. *Servlet 490 receives a response from server 460 and forwards the response to client application* (Col. 11, lines 1-67, emphasis added). Thus, “[s]ervlet 490 acts as a proxy by routing requests and responses between application 410 and serve 460” (Col. 11, lines 40-41).

Abdelnur does not teach or suggest “in response to receiving a first work request to perform first work from a first process on a first node from the plurality of nodes, determining

based upon the first work and the mapping data, that the first work is to be performed on a second node from the plurality of nodes; and providing the first work request to a second process on the second node, *wherein the first work request specifies that the first process is to directly receive results of the first work,*” as is required by Claim 1 (emphasis added). In the first scenario described above, server 460 receives a request directly from client 430, but server 460 does not provide the request to another process on another node for processing, as is required by Claim 1. Rather, server 460 itself processes the request and provides a response directly to client 430. In the second scenario, server 460 does not provide its response directly to client 430. Rather, server 460 provides its response back to Web server 480, which is then forwarded by servlet 490 to client 430. Thus, both scenarios are different than the approach recited in Claim 1.

Smith was not relied upon for the aforementioned limitations and it is respectfully submitted that *Smith* does not in any way teach or suggest the aforementioned limitations.

Claim 1 also requires “maintaining mapping data that specifies work that can be performed by each of the plurality of nodes.” The mapping data is used to determine that the first work is to be performed on the second node. The Office Action admits that *Abdelnur* does not teach or suggest this limitation, but asserts that the text at Col. 10, lines 26-42 of *Smith* does teach this limitation. This portion of *Smith* describes the use of a mapping file 110 that contains mappings of session manager IDs (MIDs) to session IDs (SIDs). Thus, the mapping file 110 indicates which session manager, if any, is handling a particular session.

It is respectfully submitted that the mapping file 110 is very different than the mapping data recited in Claim 1. The mapping data recited in Claim 1 “specifies work that can be performed by each of the plurality of nodes” so that when a work request is received, a determination can be made as to which node can process the request, based upon the work that

needs to be performed. The mapping file 110 of *Smith* only indicates which session managers are currently managing which sessions. The mapping file 110 could therefore not be used to determine which node can process a new work request for work that is different than the current active sessions. It is therefore respectfully submitted that the limitation of “maintaining mapping data that specifies work that can be performed by each of the plurality of nodes” is not taught or suggested by *Abdelnur* and *Smith*, alone or in combination.

It is also respectfully submitted that the steps in Claim 1 of “determining based upon the first work and the mapping data, that the first work is also to be performed on a third node from the plurality of nodes, and providing a second work request to a third process on the third node, wherein the second work request specifies that results of the first work performed on the third node are to be provided directly to the first process” are not in any way taught or suggested by *Guenther*. There is no teaching or suggestion in *Abdelnur* or *Smith* of providing two different work requests to two different nodes, based upon a single work request received from a process, where each of the two work requests specifies that results of performing the work are to be provided directly to the process that issued the original work request.

In view of the foregoing, it is respectfully submitted that Claim 1 contains several limitations that are not taught or suggested by *Abdelnur* and *Smith*, alone or in combination, and is therefore patentable over *Abdelnur* and *Smith*.

CLAIMS 2-4 AND 6-9

Claims 2-4 and 6-9 depend from Claim 1 and include all of the limitations of Claim 1. It is therefore respectfully submitted that Claims 2-4 and 6-9 are patentable over *Abdelnur* and *Smith* for at least the reasons set forth herein with respect to Claim 1. Furthermore, it is

respectfully submitted that Claims 2-4 and 6-9 recite additional limitations that independently render them patentable over *Abdelnur* and *Smith*.

CLAIMS 16-19

Claims 16-19 include limitations similar to Claims 1-10, except in the context of a distributed computing system. It is therefore respectfully submitted that Claims 16-19 are patentable over *Abdelnur* and *Smith* for at least the reasons set forth herein with respect to Claims 1-10.

CLAIMS 20-23 AND 25-28

Claims 20-23 and 25-28 include limitations similar to Claims 1-10, except in the context of a computer-readable medium. It is therefore respectfully submitted that Claims 20-23 and 25-28 are patentable over *Abdelnur* and *Smith* for at least the reasons set forth herein with respect to Claims 1-10.

In view of the foregoing, reconsideration and withdrawal of the rejection of Claims 1-4, 6-9, 16-23 and 25-28 under 35 U.S.C. §103(a) as being unpatentable over *Abdelnur* in view of *Smith* is respectfully requested.

REJECTION OF CLAIMS 1-29 UNDER 35 U.S.C. §102(e)

Claims 1-29 were rejected under 35 U.S.C. §102(e) as being anticipated by *Guenther et al.*, U.S. Patent No. 6,360,262 (hereinafter "*Guenther*"). This rejection is now moot with respect to canceled Claims 5, 10-15, 24 and 29. It is respectfully submitted that Claims 1-4, 6-9, 16-23 and 25-28, as amended, are not anticipated by *Guenther* for at least the reasons provided hereinafter.

CLAIM 1

Claim 1, as amended, recites a method for processing data on a distributed computing system that includes a plurality of nodes, that requires the steps of:

“maintaining mapping data that specifies work that can be performed by each of the plurality of nodes;
in response to receiving a first work request to perform first work from a first process on a first node from the plurality of nodes, determining based upon the first work and the mapping data, that the first work is to be performed on a second node from the plurality of nodes;
providing the first work request to a second process on the second node, wherein the first work request specifies that the first process is to directly receive results of the first work;
determining based upon the first work and the mapping data, that the first work is also to be performed on a third node from the plurality of nodes, and
providing a second work request to a third process on the third node, wherein the second work request specifies that results of the first work performed on the third node are to be provided directly to the first process.”

It is respectfully submitted that the steps in Claim 1 of “determining based upon the first work and the mapping data, that the first work is also to be performed on a third node from the plurality of nodes, and providing a second work request to a third process on the third node, wherein the second work request specifies that results of the first work performed on the third node are to be provided directly to the first process” are not in any way taught or suggested by *Guenther*.

In support of the rejection of Claim 1, the Office Action relied upon the description in *Guenther* of the resource router operating in “handoff” mode (Col. 4, line 64 through Col. 5, line 12 and FIG. 4B). In the “handoff” mode, the client sends a request to access a server object to the resource router. The resource router forwards the request to the “best” server, but leaves the “from” field containing the address of the client intact in the IP packets. The server sends its response directly to the client by targeting the address in the “from” field in the IP packets. This

is contrast to the proxy/gateway mode of operation, where the resource router substitutes its own address for the client address in the “from” field, so that the response from the server is sent back through the resource router.

In both the “handoff” and proxy/gateway modes of operation, the request received from the client is forwarded to the selected server, either with the contents of the “from” field left intact, or changed to the address of the resource router. There is no teaching or suggestion of providing two different work requests to two different nodes, based upon a single work request received from a process, where each of the two work requests specifies that results of performing the work are to be provided directly to the process that issued the original work request.

It is therefore respectfully submitted that the steps in Claim 1 of “determining based upon the first work and the mapping data, that the first work is also to be performed on a third node from the plurality of nodes, and providing a second work request to a third process on the third node, wherein the second work request specifies that results of the first work performed on the third node are to be provided directly to the first process” are not taught or suggested by *Guenther* and that Claim 1 is therefore not anticipated by *Guenther*.

CLAIMS 2-4 AND 6-9

Claims 2-4 and 6-9 depend from Claim 1 and include all of the limitations of Claim 1. It is therefore respectfully submitted that Claims 2-4 and 6-9 are not anticipated by *Guenther* for at least the reasons set forth herein with respect to Claim 1. Furthermore, it is respectfully submitted that Claims 2-4 and 6-9 recite additional limitations that independently render them patentable over *Guenther*.

CLAIMS 16-19

Claims 16-19 include limitations similar to Claims 1-10, except in the context of a distributed computing system. It is therefore respectfully submitted that Claims 16-19 are not anticipated by *Guenther* for at least the reasons set forth herein with respect to Claims 1-10.

CLAIMS 20-23 AND 25-28

Claims 20-23 and 25-28 include limitations similar to Claims 1-10, except in the context of a computer-readable medium. It is therefore respectfully submitted that Claims 20-23 and 25-28 are not anticipated by *Guenther* for at least the reasons set forth herein with respect to Claims 1-10.

In view of the foregoing, reconsideration and withdrawal of the rejection of Claims 1-4, 6-9, 16-23 and 25-28 under 35 U.S.C. §102(e) as being anticipated by *Guenther* is respectfully requested.

NEW CLAIMS 30-32

It is respectfully submitted that new Claims 30-32 are patentable over the cited references because none of the cited references, alone or in combination, teach or suggest “in response to receiving a first work request to perform first work from a first process on a first node from the plurality of nodes, determining based upon the first work and the mapping data, that the first work is to be performed on a second node from the plurality of nodes, generating an updated first work request that specifies that the first process is to directly receive results of performing the first work, and providing the updated first work request to a second process on the second node.” For example, the three references cited and relied upon, namely, *Abdelnur*, *Smith* and *Guenther*, all teach forwarding a request received from a client, and in the case of *Guenther*, updating the contents of a “from” field in IP packets. None of these references teach or suggest “generating an updated first work request that specifies that the first process is to directly receive results of

performing the first work.” It is therefore respectfully submitted that new Claims 30-32 are patentable over the cited references.

It is respectfully submitted that all of the pending claims are in condition for allowance and the issuance of a notice of allowance is respectfully requested. If there are any additional charges, please charge them to Deposit Account No. 50-1302.

The Examiner is invited to contact the undersigned by telephone if the Examiner believes that such contact would be helpful in furthering the prosecution of this application.

Respectfully submitted,
HICKMAN PALERMO TRUONG & BECKER LLP



Edward A. Becker
Reg. No. 37,777
Date: December 26, 2002

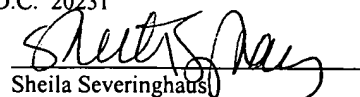
1600 Willow Street
San Jose, CA 95125
(408) 414-1204
Facsimile: (408) 414-1076

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: BOX AMENDMENT, Commissioner for Patents, Washington, D.C. 20231

on December 26, 2002

by


Sheila Severinghaus

CLAIMS IN "MARKED UP" FORM

1 1. (ONCE AMENDED) A method for processing data on a distributed computing
2 system that includes a plurality of nodes, the method comprising the steps of:
3 maintaining mapping data that specifies work that can be performed by each of the
4 plurality of nodes;
5 in response to receiving a first work request to perform first work from a first
6 process on a first node from the plurality of nodes, determining based
7 upon the first work and the mapping data, that the first work is to be
8 performed on a second node from the plurality of nodes; [and]
9 providing the first work request to a second process on the second node, wherein
10 the first work request specifies that the first process is to directly receive
11 results of the first work; [work directly from the second process.]
12 determining based upon the first work and the mapping data, that the first work is
13 also to be performed on a third node from the plurality of nodes, and
14 providing a second work request to a third process on the third node, wherein the
15 second work request specifies that results of the first work performed on
16 the third node are to be provided directly to the first process.

1 2. (NOT AMENDED) The method as recited in Claim 1, further including the steps
2 of
3 in response to receiving a second request to perform second work from the first
4 process, determining based upon the second work and the mapping data,
5 that the second work is to be performed on a third node from the plurality
6 of nodes, and

7 providing the second request to a third process on the third node, wherein the
8 second request specifies that the first process is to receive results of the
9 second work directly from the third process.

1 3. (NOT AMENDED) The method as recited in Claim 1, further including the steps
2 of
3 in response to receiving a second request to perform second work from a third
4 process on a third node from the plurality of nodes, determining based
5 upon the second work and the mapping data, that the second work is to be
6 performed on the second node, and
7 providing the second request to the second process, wherein the second request
8 specifies that the third process is to receive results of the second work
9 directly from the second process.

1 4. (NOT AMENDED) The method as recited in Claim 1, further including the steps
2 of
3 in response to receiving a second request to perform second work from a third
4 process on a third node from the plurality of nodes, determining based
5 upon the second work and the mapping data, a fourth node from the
6 plurality of nodes on which the second work is to be performed, and
7 providing the second request to a fourth process on the fourth node, wherein the
8 second request specifies that the third process is to receive results of the
9 second work directly from the fourth process.

1 5. (CANCELED) The method as recited in Claim 1, further including the steps of
2 determining based upon the first work and the mapping data, that the first work is
3 also to be performed on a third node from the plurality of nodes, and

4 providing a second request to a third process on the third node, wherein the
5 second request specifies that results of the first work performed on the
6 third node are to be provided from the third node directly to the first
7 process.

1 6. (NOT AMENDED) The method as recited in Claim 1, wherein the step of
2 determining that the first work is to be performed on a second node includes the
3 step of
4 determining one or more resources required to perform the first work, and
5 determining which of the plurality of nodes is allowed to perform the first work
6 on the one or more resources.

1 7. (NOT AMENDED) The method as recited in Claim 1, wherein:
2 the step of determining that the first work is to be performed on a second node
3 from the plurality of nodes includes the step of a director determining that
4 the first work is to be performed on a second node from the plurality of
5 nodes, and
6 the step of providing the first work request to a second process on the second
7 node includes the step of the director providing the first work request to a
8 second process on the second node.

1 8. (NOT AMENDED) The method as recited in Claim 1, further comprising the step
2 of upon completion of the first work, the second process providing the results of
3 the first work directly to the first process.

1 9. (NOT AMENDED) The method as recited in Claim 1, wherein the first work
2 request is a remote procedure call.

- 1 10. (CANCELED) The method as recited in Claim 1, further including the step of
2 generating an updated first work request that specifies that the first process is to
3 receive the first results of performing the first work and wherein the step of
4 providing the first work request to a second process includes the step of providing
5 the updated first work request to the second process.
- 1 11. (CANCELED) A method for processing data on a distributed computing system
2 that includes a plurality of nodes, the method comprising the steps of:
3 a director maintaining mapping data that specifies work that can be performed by
4 each of the plurality of nodes;
5 the director receiving a first remote procedure call from a first client process on a
6 first client node from the plurality of nodes, wherein the first remote
7 procedure call requests that first work be performed and that results of the
8 first work be provided directly to the first client process;
9 the director examining the first remote procedure call and determining based upon
10 the first work and the mapping data, that a first server process on a first
11 server node from the plurality of nodes is to perform the first work; and
12 the director providing the first remote procedure call to the first server process.
- 1 12. (CANCELED) The method as recited in Claim 11, further including the steps of
2 the director determining one or more resources required to perform the first work,
3 and
4 the director determining which of the plurality of nodes have permission to
5 perform a requested operation on the one or more resources required to
6 perform the first work.

- 1 13. (CANCELED) The method as recited in Claim 12, wherein the step of the
2 director determining the resources available on the plurality of nodes includes the
3 director examining resource data associated with the plurality of nodes.
- 1 14. (CANCELED) The method as recited in Claim 11, further including the steps of
2 the director receiving a second remote procedure call from the first client process,
3 wherein the second remote procedure call requests that second work be
4 performed and that results of the second work be provided directly to the
5 first client process,
6 the director examining the second remote procedure call and determining based
7 upon the second work and the mapping data, that a second server process
8 on a second server node from the plurality of nodes is to perform the
9 second work, and
10 the director providing the second remote procedure call to the second server
11 process.
- 1 15. (CANCELED) The method as recited in Claim 11, further including the steps of
2 the director receiving a second remote procedure call from a second client process
3 on a second client node in the plurality of nodes, wherein the second
4 remote procedure call requests that second work be performed and that
5 results of the second work be provided directly to the second client
6 process,
7 the director examining the second remote procedure call and determining based
8 upon the second work and the mapping data, that a second server process
9 on a second server node from the plurality of nodes is to perform the
10 second work, and
11 the director providing the second remote procedure call to the second server
12 process.

1 16. (ONCE AMENDED) A distributed computing system for performing work, the
2 distributed computing system comprising:
3 a plurality of nodes; and
4 a director communicatively coupled to the plurality of nodes, the director being
5 configured to
6 maintain mapping data that specifies work that can be performed by each
7 of the plurality of nodes,
8 in response to a first work request to perform first work from a first
9 process on a first node from the plurality of nodes, determine based
10 upon the first work and the mapping data, that the first work is to
11 be performed on a second node from the plurality of nodes, and
12 request that the first work be performed by a second process on the second
13 node, wherein the request specifies that first results of the first
14 work be directly provided [from the second process directly] to the
15 first [process.] process;
16 determining based upon the first work and the mapping data, that the first
17 work is also to be performed on a third node from the plurality of
18 nodes, and
19 providing a second work request to a third process on the third node,
20 wherein the second work request specifies that results of the first
21 work performed on the third node are to be provided directly from
22 the third node to the first process.

1 17. (NOT AMENDED) The distributed computing system as recited in Claim 16,
2 wherein the director is further configured to provide the first work request to the
3 second process.

- 1 18. (NOT AMENDED) The distributed computing system as recited in Claim 16,
2 wherein the director is further configured to
3 generate a second work request to request that the second process perform the
4 first work and provide the first results directly to the first process, and
5 provide the second work request to the second process.
- 1 19. (NOT AMENDED) The distributed computing system as recited in Claim 16,
2 further comprising resource data that specifies the access rights of the plurality of
3 nodes relative to resources.
- 1 20. (ONCE AMENDED) A computer-readable medium carrying one or more
2 sequences of one or more instructions for processing data on a distributed
3 computing system that includes a plurality of nodes, the one or more sequences of
4 one or more instructions include instructions which, when executed by one or
5 more processors, cause the one or more processors to perform the steps of:
6 maintaining mapping data that specifies work that can be performed by each of
7 the plurality of nodes;
8 in response to receiving a first work request to perform first work from a first
9 process on a first node from the plurality of nodes, determining based
10 upon the first work and the mapping data, that the first work is to be
11 performed on a second node from the plurality of nodes; [and]
12 providing the first work request to a second process on the second node, wherein
13 the first work request specifies that the first process is to directly receive
14 results of the first work; [work directly from the second process.]
15 determining based upon the first work and the mapping data, that the first work is
16 also to be performed on a third node from the plurality of nodes, and

17 providing a second work request to a third process on the third node, wherein the
18 second work request specifies that results of the first work performed on
19 the third node are to be provided directly from the third node to the first
20 process.

1 21. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 further including the steps of
3 in response to receiving a second request to perform second work from the first
4 process, determining based upon the second work and the mapping data,
5 that the second work is to be performed on a third node from the plurality
6 of nodes, and
7 providing the second request to a third process on the third node, wherein the
8 second request specifies that the first process is to receive results of the
9 second work directly from the third process.

1 22. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 further including the steps of
3 in response to receiving a second request to perform second work from a third
4 process on a third node from the plurality of nodes, determining based
5 upon the second work and the mapping data, that the second work is to be
6 performed on the second node, and
7 providing the second request to the second process, wherein the second request
8 specifies that the third process is to receive results of the second work
9 directly from the second process.

1 23. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 further including the steps of

3 in response to receiving a second request to perform second work from a third
4 process on a third node from the plurality of nodes, determining based
5 upon the second work and the mapping data, a fourth node from the
6 plurality of nodes on which the second work is to be performed, and
7 providing the second request to a fourth process on the fourth node, wherein the
8 second request specifies that the third process is to receive results of the
9 second work directly from the fourth process.

1 24. (CANCELED) The computer-readable medium as recited in Claim 20, further
2 including the steps of
3 determining based upon the first work and the mapping data, that the first work is
4 also to be performed on a third node from the plurality of nodes, and
5 providing a second request to a third process on the third node, wherein the
6 second request specifies that results of the first work performed on the
7 third node are to be provided from the third node directly to the first
8 process.

1 25. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 wherein the step of determining that the first work is to be performed on a second
3 node includes the step of
4 determining one or more resources required to perform the first work, and
5 determining which of the plurality of nodes is allowed to perform the first work
6 on the one or more resources.

1 26. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 wherein:
3 the step of determining that the first work is to be performed on a second node
4 from the plurality of nodes includes the step of a director determining that

5 the first work is to be performed on a second node from the plurality of
6 nodes, and
7 the step of providing the first work request to a second process on the second
8 node includes the step of the director providing the first work request to a
9 second process on the second node.

1 27. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 further comprising the step of upon completion of the first work, the second
3 process providing the results of the first work directly to the first process.

1 28. (NOT AMENDED) The computer-readable medium as recited in Claim 20,
2 wherein the first work request is a remote procedure call.

1 29. (CANCELED) The computer-readable medium as recited in Claim 20, further
2 including the step of generating an updated first work request that specifies that
3 the first process is to receive the first results of performing the first work and
4 wherein the step of providing the first work request to a second process includes
5 the step of providing the updated first work request to the second process.

1 30. (NEW) A method for processing data on a distributed computing system that
2 includes a plurality of nodes, the method comprising the steps of:
3 maintaining mapping data that specifies work that can be performed by each of the
4 plurality of nodes; and
5 in response to receiving a first work request to perform first work from a first
6 process on a first node from the plurality of nodes,
7 determining based upon the first work and the mapping data, that the first
8 work is to be performed on a second node from the plurality of
9 nodes,

10 generating an updated first work request that specifies that the first process
11 is to directly receive results of performing the first work, and
12 providing the updated first work request to a second process on the second
13 node.

1 31. (NEW) An apparatus for processing data on a distributed computing system, the
2 apparatus comprising a memory carrying one or more sequences of one or more
3 instructions which, when executed by one or more processors, cause the one or more
4 processors to perform the steps of:
5 maintaining mapping data that specifies work that can be performed by each of the
6 plurality of nodes; and
7 in response to receiving a first work request to perform first work from a first
8 process on a first node from the plurality of nodes,
9 determining based upon the first work and the mapping data, that the first
10 work is to be performed on a second node from the plurality of
11 nodes,
12 generating an updated first work request that specifies that the first process
13 is to directly receive results of performing the first work, and
14 providing the updated first work request to a second process on the second
15 node.

1 32. (NEW) A computer-readable medium for processing data on a distributed
2 computing system, the computer-readable medium carrying one or more sequences
3 of one or more instructions which, when executed by one or more processors, cause
4 the one or more processors to perform the steps of:
5 maintaining mapping data that specifies work that can be performed by each of the
6 plurality of nodes; and

7 in response to receiving a first work request to perform first work from a first
8 process on a first node from the plurality of nodes,
9 determining based upon the first work and the mapping data, that the first
10 work is to be performed on a second node from the plurality of
11 nodes,
12 generating an updated first work request that specifies that the first process
13 is to directly receive results of performing the first work, and
14 providing the updated first work request to a second process on the second
15 node.